

CLAIMS

What is claimed is:

- 1 1. An apparatus for acknowledging a data transfer, comprising:
2 a first protocol that is adapted to generate a request for a data transfer; and
3 a second protocol that is adapted to:
4 receive the request for the data transfer from the first protocol;
5 determine whether the request for the data transfer contains a request for
6 acknowledgement of completion of the data transfer;
7 send a performance request corresponding to the request for data transfer
8 to a third protocol; and
9 if the request for data transfer does contain a request for
10 acknowledgement of the completion of the data transfer, set a variable in
11 memory to wait for an event to correspond to the completion the request for data
12 transfer and send an acknowledgement to the first protocol upon the occurrence
13 of the event.

- 1 2. The apparatus set forth in claim 1, wherein the first protocol is an internet
2 small computer systems interface ("iSCSI") protocol.

1 3. The apparatus set forth in claim 1, wherein the second protocol is an
2 internet small computer systems interface extensions for remote direct memory access
3 (“iSER”) protocol.

1 4. The apparatus set forth in claim 1, wherein the request for the data
2 transfer comprises an attribute that indicates the request for acknowledgement of
3 completion of the data transfer.

1 5. The apparatus set forth in claim 4, wherein a value of an error recovery
2 level is notified to the second protocol from the first protocol.

1 6. The apparatus set forth in claim 1, wherein the third protocol is a remote
2 direct memory access (“RDMA”) protocol.

1 7. The apparatus set forth in claim 1, wherein the event relates to a zero
2 length remote direct memory access (“RDMA”) read completion.

1 8. A network, comprising:
2 a plurality of systems, at least one of the plurality of systems comprising a
3 protocol stack and a process;
4 at least one input/output device;

5 a network that connects the plurality of systems and the at least one input/output
6 device for communication; and
7 wherein the protocol stack comprises:
8 a first protocol layer that interacts with the consumer;
9 a second protocol layer that is adapted to:
10 receive a data exchange request from the first protocol layer;
11 examine the data exchange request to determine if an
12 acknowledgement request is indicated;
13 send a performance request corresponding to the data exchange
14 request to a third protocol layer; and
15 if the data exchange request contains the acknowledgement
16 request, set a variable in memory to wait for an event that corresponds to
17 the completion of the performance request and send an acknowledgement
18 to the first protocol layer upon the occurrence of the event.

1 9. The network set forth in claim 8, wherein the third protocol layer
2 interacts with the second protocol layer and is adapted to:
3 receive the performance request that corresponds to the data exchange
4 request; and
5 transmit a message to one of the at least one of the plurality of systems
6 and the at least one input/output device via the network.

1 10. The network set forth in claim 9, comprising a remote direct memory
2 access network interface card ("RNIC") that is used by the protocol stack to exchange
3 the message between the at least one of the plurality of systems and the at least one
4 input/output device via the network.

1 11. The network set forth in claim 9, wherein the message is a remote direct
2 memory access ("RDMA") write message.

1 12. The network set forth in claim 9, wherein the message is a zero length
2 remote direct memory access ("RDMA") read message.

1 13. The network set forth in claim 8, wherein the second protocol layer is an
2 internet small computer systems interface extensions for remote direct memory access
3 ("iSER") protocol.

1 14. The network set forth in claim 8, wherein the data exchange request
2 comprises an attribute and data.

1 15. The network set forth in claim 8, wherein the process is a small computer
2 systems interface protocol ("SCSI").

1 16. A method of acknowledging a data transfer, the method comprising:

2 receiving a request for a data transfer from a first protocol;
3 determining whether the request for the data transfer contains a request for
4 acknowledgement of completion of the data transfer;
5 sending a performance request corresponding to the request for data transfer to a
6 second protocol; and
7 if the request for data transfer does contain a request for acknowledgement of
8 completion of the data transfer, setting a variable in memory to wait for
9 an event corresponding to completion of the data transfer and sending an
10 acknowledgement to the first protocol upon the occurrence of the event.

1 17. The method set forth in claim 16, comprising defining the first protocol
2 as an internet small computer systems interface ("iSCSI") protocol.

1 18. The method set forth in claim 16, comprising defining the second
2 protocol as a remote direct memory access ("RDMA") protocol.

1 19. The method set forth in claim 16, comprising defining the event to relate
2 to a zero length remote direct memory access ("RDMA") read message completion.

1 20. The method set forth in claim 16, comprising defining the event to relate
2 to a remote direct memory access ("RDMA") write message completion.

1 21. The method set forth in claim 16, comprising establishing an error
2 recovery level by the first protocol to indicate the error recovery level in the request for
3 acknowledgement of completion of the data transfer.

1 22. An apparatus for acknowledging a data transfer, comprising:
2 means for receiving a request for a data transfer from a first protocol;
3 means for determining whether the request for the data transfer contains a
4 request for acknowledgement of completion of the data transfer in a second protocol;
5 means for sending a performance request corresponding to the request for data
6 transfer to a third protocol; and
7 means for setting a variable in memory to wait for an event to correspond to the
8 completion of the performance request and sending an acknowledgement to the first
9 protocol upon the occurrence of the event if the request for the data transfer does
10 contain the request for acknowledgement of completion of the data transfer.

1 23. A program for acknowledging a data transfer, comprising:
2 a machine readable medium;
3 a first protocol stored on the machine readable medium, the first protocol being
4 adapted to generate a request for a data transfer; and
5 a second protocol stored on the machine readable medium, the second protocol
6 being adapted to:
7 receive the request for the data transfer from the first protocol;

8 determine whether the request for the data transfer contains a request for
9 acknowledgement of completion of the data transfer;
10 send a performance request corresponding to the request for data transfer
11 to a third protocol; and
12 set a variable in memory to wait for an event to correspond to the
13 completion of the performance request and send an acknowledgement to the first
14 protocol upon the occurrence of the event if the request for data transfer does
15 contain a request for acknowledgement of completion of the data transfer.